

CLAIMS

What is claimed is:

1. A switch in a network being logically represented by a spanning tree with cross-links, wherein the spanning tree identifies a root switch and a plurality of other switches, the switch comprising:
 - means for receiving at a current switch a copy of a broadcast packet of a broadcast from each switch for which the current switch is a descendent switch;
 - means for selecting a copy of the broadcast packet of the broadcast; and
 - means for forwarding copies of the selected copy of the broadcast packet of the broadcast to all descendents of the current switch.
2. A switch as recited in claim 1, wherein the means for receiving includes a plurality of switch ports, each port being configured to connect to a link in the network.
3. A switch as recited in claim 1, wherein the means for selecting includes a routing control circuit that control the routing of packets in the switch.
4. A switch as recited in claim 1,
 - wherein the means for receiving includes a plurality of switch ports, each port being configured to connect to a link in the network; and
 - wherein the means for forwarding includes a crossbar connection unit that is configured to interconnect the switch ports.
5. A switch for a mesh-connected network, that is logically represented by a spanning tree with cross-links, the spanning tree identifying therein a root switch and a plurality of other switches, the switch comprising:
 - a plurality of switch ports that are each configured to connect to a link in the network and to receive, in a broadcast, from a connected link, a copy of a broadcast packet from each switch for which the switch is a descendent switch;

a crossbar connection unit that is connected to the routing logic circuit and the plurality of switch ports, and configured, in a broadcast, to forward, from a selected switch port, copies of the broadcast packet of the broadcast to any ports having connected links to descendents of the switch in the network; and

a routing logic circuit that is connected to the crossbar connection unit to control the routing of packets in the crossbar connection unit, and configured, in a broadcast, to select one of the switch ports that has a copy of the broadcast packet of the broadcast and any other ports as output ports for transmitting the selected copy of the broadcast packet to descendents of the switch.

6. A switch as recited in claim 5, wherein the crossbar connection unit is configured to forward copies of the selected copy of the broadcast packet simultaneously.

7. A switch as recited in claim 5,

wherein a link connected to the switch is a cross-link to a switch for which the switch is a descendent switch; and

wherein a copy of the broadcast packet sent on the cross-link is represented by a token that identifies the broadcast packet.

8. A switch as recited in claim 7, wherein the token includes header information that identifies the broadcast packet.

9. A switch as recited in claim 5, wherein the crossbar unit is configured, prior to forwarding copies of the selected copy of the broadcast packet, to retain, at the current switch, the selected copy for forwarding; and to discard, at the current switch, all other copies of the broadcast packet.

10. A switch as recited in claim 5,

wherein a port of the switch receives, in a broadcast, a copy of the broadcast packet from at least one switch for which the switch is a descendent switch after a predetermined amount of time; and

further comprising a switch control processor that is connected to the crossbar connection unit and the routing logic circuit, and configured to detect that a failure occurs if copy of a

broadcast packet is received after the predetermined amount of time and to initialize the switch in response thereto.

11. A switch as recited in claim 10, wherein the switch control processor is configured to initialize the switch by:

- propagating an initialization state to any other switches that are connected to the switch;
- collecting network topology information from any other switches that are connected to the switch;
- distributing the network topology information to the other switches; and
- waiting until all switches in the network are initialized before resuming normal operation.

12. A switch as recited in claim 5, wherein the routing logic unit is configured to recognize at least two concurrent but separate broadcasts in the network.

13. A switch as recited in claim 12,

wherein each separate broadcast has an identifier in each broadcast packet associated with the broadcast; and

wherein the routing logic unit is further configured to recognize broadcast packets as belonging to one of the at least two broadcasts by being further configured to examine the identifier in the broadcast packet.

14. A method for broadcasting packets as recited in claim 13, wherein the identifier incorporated into each copy of the broadcast packets for each identified broadcast is a sequence number.

15. A method for broadcasting packets as recited in claim 13, wherein the identifier in each broadcast is unique to the broadcast.

16. A switch as recited in claim 12,

further comprising a switch control processor that is connected to the crossbar connection unit and the routing logic unit, and configured to initiate a broadcast in response to a request from a host connected to the switch; and

wherein the switch control processor is further configured to include a unique identifier in a broadcast packet to distinguish the broadcast from the at least one other broadcast in the network.

17. A switch as recited in claim 5, wherein the crossbar connection unit is further configured to start the forwarding of copies of the selected copy of the broadcast packet of the broadcast to all descendents of the switch before storing a complete copy of the broadcast packet from each from each switch for which the switch is a descendent switch.

18. A switch as recited in claim 5, wherein the crossbar connection unit and routing logic circuit are each further configured to forward point-to-point packets in the switch concurrently with forwarding broadcast packets in the switch.

19. A mesh connected network for interconnecting a plurality of hosts, the network comprising:
a plurality of switches for routing data packets between hosts in the network; and
a plurality of point-to-point links, each providing a connection between a pair of switches or between a switch and a host in the network;

wherein each switch includes:

a plurality of switch ports that are each configured to connect to a link in the network and to receive, in a broadcast, from a connected link, a copy of a broadcast packet from each switch for which the switch is a descendent switch;

a crossbar connection unit that is connected to the routing logic circuit and the plurality of switch ports, and configured, in a broadcast, to forward, from a selected switch port, copies of the broadcast packet of the broadcast to any ports having connected links to descendents of the switch in the network; and

a routing logic circuit that is connected to the crossbar connection unit to control the routing of packets in the crossbar connection unit, and configured, in a broadcast, to select one of the switch ports that has a copy of the broadcast packet of the broadcast and any other ports as output ports for transmitting the selected copy of the broadcast packet to descendents of the switch.